

Deploying Firewalls

William Fithen

Julia Allen

Ed Stoner

May 1999

SECURITY IMPROVEMENT MODULE
CMU/SEI-SIM-008



Carnegie Mellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

Deploying Firewalls

CMU/SEI-SIM-008

William Fithen

Julia Allen

Ed Stoner

May 1999

Networked Systems Survivability Program

Unlimited distribution subject to the copyright.

This work is sponsored by the USAF Embedded Computer Resources Support Improvement Program (ESIP).

The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright © 1999 by Carnegie Mellon University.

Requests for permission to reproduce this document or to prepare derivative works of this document should be addressed to the SEI Licensing Agent.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number F19628-95-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 52.227-7013.

For informatin about how to order paper copies of this document, plese visit the Publications portion of the SEI Web site, <http://www.sei.cmu.edu/publications/pubweb.html>

Table of Contents

Preface	iii
Deploying Firewalls	1
1. Design the firewall system.	7
2. Acquire firewall hardware and software.	23
3. Acquire firewall documentation, training, and support.	27
4. Install firewall hardware and software.	29
5. Configure IP routing.	33
6. Configure firewall packet filtering.	35
7. Configure firewall logging and alert mechanisms.	41
8. Test the firewall system.	45
9. Install the firewall system.	55
10. Phase the firewall system into operation.	57

Preface

This document is one of a series of publications of the Software Engineering Institute at Carnegie Mellon University called *security improvement modules*. They are intended to provide practical guidance to help organizations improve the security of their networked computer systems.

Module structure

Each module addresses an important but relatively narrowly defined problem in network and system security. The first section of the module describes the problem and outlines a set of *security improvement practices* to help solve it. Each practice is a recommended way of performing common tasks related to the secure operation of networked computer systems.

The remaining sections of the module are detailed descriptions of the practices. Each includes a rationale for the recommended actions and a description of how to perform them.

Intended audience

The practices are primarily written for system and network administrators whose day-to-day activities include installation, configuration, and maintenance of the computers and networks. Occasionally, practices are written to assist the managers responsible for network and system administration.

Revised versions

Network and system technologies continue to evolve rapidly, leading to new security problems and solutions. Modules and practices need to be revised occasionally, so to permit more timely publication of new versions, we also publish them on the World Wide Web. At the end of each section of this document is the URL of its Web version.

Implementation details

How an organization adopts and implements the practices often depends on the networking and computing technologies it uses. For some practices, technology-specific implementation details are published on the World Wide Web. The Web version of each practice contains links to the implementation details.

Acknowledgements

This report and the effort to produce it were sponsored by the USAF Embedded Computer Resources Support Improvement Program (ESIP). The authors are pleased to acknowledge LtCol Joseph Jarzombek for engaging as reviewers and collaborators the USAF Information Warfare Center (AFIWC) under Mr. Feliciano Rodriguez, Director, Engineering and Analysis. AFIWC assisted in the selection of the Security Improvement Module title and content. AFIWC provided expert review and recommendations. The authors believe that this collaboration resulted in a better module for ESIP, AFIWC, and the community as a whole.

The authors appreciate the support and cooperation of all AFIWC personnel that contributed to *Deploying Firewalls*. We would like to specifically recognize the individuals that we interfaced with: Feliciano Rodriguez, Jose Linero, Joe Cano, James Dennis, Lt Lynn Blankenship, and Lt Paul Townley.

The authors would like to acknowledge the contributions of the following external reviewers whose comments greatly enhanced the quality of this document:

AFIWC:

Mr. Jose Linero
Lt Paul Townley

AUSCERT:

Rob McMillan

SEI:

Greg Gravenstreter
Jeff Havrilla
Eric J. Hayes
Steve Kalinowski
Klaus-Peter Kossakowski
Marty Lindner
Rudy Maceyko
Derek Simmel

Deploying Firewalls

A firewall is a combination of hardware and software used to implement a security policy governing the network traffic between two or more networks, some of which may be under your administrative control (e.g., your organization's networks) and some of which may be out of your control (e.g., the Internet). A network firewall commonly serves as a primary line of defense against external threats to your organization's computer systems, networks, and critical information. Firewalls can also be used to partition your organization's internal networks, reducing your risk from insider attacks.

Firewall technologies have entered into the mainstream. The "1999 Computer Security Institute/FBI Computer Crime and Security Survey" [Power 99] indicates that 91 percent of the organizations surveyed already deploy firewalls. Articles and other references covering evaluation, selection, and configuration of firewall technologies are now common in the popular press (see References at the end of this section).

However, there has been little published about designing, installing, deploying, operating, and maintaining firewalls. The practices in this module will address designing, installing, and deploying firewalls.

The term firewall is taken from the structural analog whose purpose is to slow the spread of fire in a building. In the computer literature, popular press, and vendor marketing materials, the term is used in many ways. Some people use it to identify a specific hardware component or software package, while others consider the entire collection of systems and software deployed between two networks to be parts of a firewall.

Throughout these practices, we will generally use the term firewall as an adjective modifying a noun (such as system, hardware, software, product) to make the reference clear. When we use the term firewall as a noun, we mean the general concept of a technological mechanism for the enforcement of a network traffic security policy. While this may seem cumbersome at times, we believe these distinctions will increase your understanding of our intent.

Who should read these practices

These practices are intended primarily for experienced system and network administrators and integrators.

These practices are applicable to your organization if its information infrastructure either includes or will soon include

- interconnections between internal networks and networks not under its administrative control, such as the Internet or business partner networks
- interconnections among internal networks with different security requirements

The purpose of this module is to cover the fundamentals of firewall functionality (packet filtering) and the deployment process. These practices assume that your desired firewall architecture includes packet filtering as a first step. Later versions of this module will address additional firewall capabilities such as proxies and VPNs (virtual private networks).

What these practices do not cover

These practices do not address

- the creation of a detailed security policy including the policy to be enforced by the firewall
- the evaluation and selection of specific firewall products
- post-deployment operation and maintenance of firewalls
- the design and deployment of more advanced firewall capabilities, such as
 - proxies (including SOCKS)
 - stateful inspection or dynamic packet filtering
 - network address translation
 - virtual private networks
 - Internet Protocol version 6 or other non-Internet Protocol version 4 protocols
 - network and host intrusion detection technologies
- networking fundamentals, such as
 - specific Internet protocols
 - routing and route management
 - switching and VLANs (virtual local area networks)
- system management fundamentals, such as
 - operating systems installation and maintenance
 - application software installation and maintenance
 - host intrusion detection technologies
- cryptography and encryption technologies

Security issues

Increasingly, organizations are connecting to the Internet to establish a business and electronic commerce presence and to access information rapidly. When your organization's networks are connected to the Internet without adequate security measures in place, you become vulnerable to attacks from external adversaries. Without firewalls, you will be unable to prevent many forms of undesirable access to your networks, systems, and information assets. The risks include

- loss of confidentiality of business information (e.g., financial records, strategic planning data, engineering models and prototypes, marketing plans, medical records, as well as inability to guarantee the integrity of such information)

- loss of availability of mission-critical services such as EDI (electronic data interchange), ERP (enterprise resource planning), just-in-time inventory controls, and electronic mail
- exposure of critical data about your information infrastructure that can be used by your adversaries in planning their attacks
- legal liability, regulatory liability, or public loss of confidence when your adversaries use one of your computers to carry out attacks against other organizations
- vandalism of public information services (such as your public Web site)

The use of firewall technology provides you with one of the most effective tools available to manage your networks' risk by providing you with access control mechanisms that can implement complex security policies.

Security improvement approach

To effectively deploy firewall technology, we recommend a four-part approach. It requires implementing security practices in these areas:

- preparing for firewall system deployment
- configuring your firewall system to reflect your security policy
- testing your firewall system to ensure it performs according to your specifications
- deploying the correctly configured firewall system

Summary of recommended practices

Area	Recommended Practice
Prepare	1. Design the firewall system.
Configure	2. Acquire firewall hardware and software. 3. Acquire firewall documentation, training, and support. 4. Install firewall hardware and software. 5. Configure IP routing. 6. Configure firewall packet filtering. 7. Configure firewall logging and alert mechanisms.
Test	8. Test the firewall system.
Deploy	9. Install the firewall system. 10. Phase the firewall system into operation.

Abbreviations used in these practices

DG	default gateway
DHCP	dynamic host configuration protocol
DMZ	demilitarized zone
DNS	domain name service
EDI	electronic data interchange
ERP	enterprise resource planning
FTP	file transfer protocol
HTTP	hypertext transfer protocol
ICMP	Internet control message protocol
IDS	intrusion detection system
IP	Internet protocol

ISP	Internet service provider
LDAP	lightweight directory access protocol
NAT	network address translation
NFS	network file system
NTP	network time protocol
OS	operating system
OSPF	open shortest path first
RAM	random access memory
RCS	revision control system
RIP	routing information protocol
SCCS	software configuration control system
SOCKS	general purpose application proxy ¹
SMTP	simple mail transfer protocol
SNMP	simple network management protocol
SPAK	Send PAcKets ²
SSH	secure shell
SSL	secure socket layer
TCP	transmission control protocol
UDP	user datagram protocol
VLAN	virtual local area network
VPN	virtual private network
WWW	World Wide Web

References

General firewall references:

- [Cheswick 94] Cheswick, William R. & Bellovin, Steven M. *Firewalls and Internet Security*. Reading, MA: Addison-Wesley, 1994.
- [Chapman 95] Chapman, D. Brent & Zwicky, Elizabeth D. *Building Internet Firewalls*. Sebastopol, CA: O'Reilly & Associates, 1995.
- [Cooper 97] Cooper, Deborah & Pfleeger, Charles. "Firewalls: An Expert Roundtable." *IEEE Software*, New York, NY: IEEE, September/October 1997.
- [Goncalves 98] Goncalves, Marcus. *Firewalls Complete*. New York, NY: McGraw Hill, 1998.
- [Hall 96] Hall, Eric. "Internet Firewall Essentials." Network Computing Online. Manhasset, NY: CMP Media, Inc., November, 1996. Available at <http://www.networkcomputing.com/netdesign/wall1.html>.
- [ICSA 98] *Third Annual Firewall Industry Guide*. International Computer Security Association, 1998. Available at <http://www.icsa.net/fwcd/fwcdindex.shtml>.
- [Lodin 98] Lodin, Steve & Schuba, Christoph. "Firewalls fend off invasions from the Net." *IEEE Spectrum*. New York, NY: IEEE, February, 1998.

1. Refer to <http://www.socks.nec.com>

2. A network traffic generator tool available at the COAST web site

- [Luk 98] Luk, Ellis, et al. *Protect and Survive: Using IBM Firewall 3.1 for AIX, 3rd edition*. Research Triangle Park, NC: IBM, 1998. Available at <http://www.redbooks.ibm.com>.
- [NIST 98] *Internet Security Policy: A Technical Guide*. Washington, DC: National Institute of Standards and Technology, 1998. Available at <http://csrc.nist.gov/isptg>.
- [Power 99] Power, Richard. "1999 CSI/FBI Computer Crime and Security Survey." *Computer Security Journal, Volume XV, Number 2*. San Francisco, CA: Computer Security Institute, 1999.
- [SC 99] "Firewalls Market Survey." *SC Magazine*. Framingham, MA: West Coast Publishing, Inc., April, 1999. Available at <http://www.infosecnews.com>.

Specific firewall technologies:

- [Avolio 98] Avolio, Blask. "Application Gateways and Stateful Inspection: A Brief Note Comparing and Contrasting." Trusted Information Systems, Inc., 1998. Available at <http://www.avolio.com/apgw+spf.html>.
- [Check Point 98] "Stateful Inspection Firewall Technology Tech Note." Check Point Software Technologies Ltd., 1998. Available at <http://www.checkpoint.com/products/technology/stateful1.html>.

Protocols:

- [Comer 95] Comer, Douglas E. *Internetworking with TCP/IP, volume 1: principles, protocols, and architecture. 3rd Edition*. New York, NY: Prentice-Hall, 1995.
- [Stevens 94] Stevens, W. Richard. *TCP/IP Illustrated, Volume 1: The Protocols*. Reading, MA: Addison-Wesley, 1994.

Detecting intrusions:

- [Escamilla 98] Escamilla, Terry. *Intrusion Detection: Network Security Beyond the Firewall*. New York, NY: Wiley Computer Publishing, 1998.
- [Firth 97] Firth, Robert, et al. *Detecting Signs of Intrusion*. (CMU/SEI-SIM-001). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1997. Available at <http://www.cert.org/security-improvement/m01.html>.
- [Kochmar 98] Kochmar, John, et al. *Preparing to Detect Signs of Intrusion*. (CMU/SEI-SIM-005). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1998. Available at <http://www.cert.org/security-improvement/m05.html>.

Architecture tradeoff analysis:

- [Kazman 98] Kazman, Rick, et al. "The Architecture Tradeoff Analysis Method." *Proceedings of the Fourth IEEE International Conference on Engineering of Complex Computer Systems (ICECCS)*. Monterey, CA: IEEE, August 1998, 68-78.
- [Kazman 99] Kazman, Rick, et al. "Experience with Performing Architecture Tradeoff Analysis." *Proceedings of ICSE 21*. Los Angeles, CA: ACM Press, May 1999, soon to be published.

Linux systems:

- [RedHat] Red Hat Software. Available at <http://www.redhat.com>.
- [FreshMeat] Lenz, Patrick; Edmonds, Robert; Thompson, Christoph; & Weaver, Ryan. Available at <http://freshmeat.net>.
- [Grennan 96] Grennan, Mark. *Firewalling and Proxy Server HOWTO*. Version 0.4. November 8, 1996. Available at <http://metalab.unc.edu/LDP/HOWTO/Firewall-HOWTO.html>.
- [Russell 98] Russell, Paul. *Linux IPCHAINS-HOWTO*. Version 1.0.5. October 27, 1998. Available at <http://metalab.unc.edu/LDP/HOWTO/IPCHAINS-HOWTO.html>.

Other related technologies:

- [Schneier 96] Schneier, Bruce. *Applied Cryptography*, 2nd Edition. New York, NY: John Wiley & Sons, Inc., 1996.
- [Simmel 99] Simmel, Derek, et al. *Securing Desktop Workstations*. (CMU/SEI-SIM-004). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1999. Available at <http://www.cert.org/security-improvement/m04.html>.

Other information

Information on firewall issues can be found at the mailing list archive maintained by Gnac at <http://lists.gnac.net/firewalls/>. This site includes a link to the Internet Firewalls FAQ (frequently asked questions).

Check the COAST (Computer Operations, Audit, and Security Technology) website at Purdue University. Firewall-related materials can be found at <http://www.cs.purdue.edu/coast/firewalls/fw-body.html>. The site contains references for relevant books, papers, articles, reports, guides, research, products, firewall testing results, firewall tools, network tools, system monitoring tools, mailing lists, newsgroups, conferences, and frequently asked questions.

Where to find updates

The latest version of this module is available on the Web at URL <http://www.cert.org/security-improvement/modules/m08.html>

1

Design the firewall system.

Designing a firewall requires that you understand and identify the boundaries between security domains in your network. A network security domain is a contiguous region of a network that operates under a single, uniform security policy. Wherever these domains intersect, there is a potential need for a policy conflict resolution mechanism at that boundary. This is where firewall technology can help.

The most common boundary where firewalls are applied today is between an organization's internal networks and the Internet. When establishing an Internet firewall, the first thing you must decide is its basic architecture (assuming you have previously established your firewall requirements¹ and the security policy it is intended to implement). In this context, architecture refers to the inventory of components (hardware and software), and the connectivity and distribution of functions among them. There are two classes of firewall architectures, which we refer to as the *single layer* and the *multiple layer* architectures.

In a single layer architecture (see figure 1-1 at the end of this section), one network host is allocated all firewall functions and is connected to each network for which it is to control access. This approach is usually chosen when containing cost is a primary factor or when there are only two networks to interconnect. It has the advantage that everything there is to know about the firewall resides on that one host. In cases where the policy to be implemented is simple and there are few networks being interconnected, this approach can also be very cost-effective to operate and maintain over time. The greatest disadvantage of the single layer approach is its susceptibility to implementation flaws or configuration errors — depending on the type, a single flaw or error might allow firewall penetration.

-
1. These should have been specified during your firewall evaluation and selection process. Areas you should have considered include
 - risks you are trying to mitigate with the firewall (i.e., the information assets and resources you are trying to protect and the threats that you are trying to protect against)
 - services you intend to offer to the Internet from your network
 - services you intend to use on the Internet from your network
 - identification of the users of these services
 - firewall availability and performance requirements
 - determining who will manage the firewall system and how they will manage it
 - determining the system and network growth that the firewall system will need to accommodate in the future

Other considerations can be found in *Firewalls Complete* [Goncalves 98] and the *Third Annual Firewall Industry Guide* [ICSA 98].

In a multiple layer architecture (see figure 1-2 at the end of this section), the firewall functions are distributed among a small number of hosts, typically connected in series, with DMZ networks between them. This approach is more difficult to design and operate, but can provide substantially greater security by diversifying the defenses you are implementing. Although more costly, we advise using different technology in each of these firewall hosts. This reduces the risk that the same implementation flaws or configuration errors will exist in every layer. The most common design approach for this type of architecture is an Internet firewall composed of two hosts interconnected with one DMZ network.

Having chosen the basic architecture (i.e., the number of hosts, the method in which they are connected, the tasks that each will perform), the next step is to select the firewall functions to be implemented in these hosts. The two most basic categories of firewall function are packet filtering and application proxies. These functions can be used separately or jointly and can be implemented on the same or on different firewall hosts. Recently, packet filtering firewall products have gained some of the features of application proxies and are generally referred to as stateful inspection packet filters. See *Building Internet Firewalls* [Chapman 95], *Firewalls Complete* [Goncalves 98], and “Firewalls fend off invasions from the Net.” [Lodin 98] for a more detailed explanation of the different types of firewall functions.

There are good reasons to use both packet filtering and application proxies. Certain services (e.g., SMTP, HTTP, or NTP) are usually safe to control via packet filters while others (e.g., DNS, FTP) may require the more complex features available only in proxies. Packet filtering is fast, while application proxies are generally slower. In cases where greater access control is required and the poorer performance of proxies cannot be tolerated, stateful inspection packet filters may be an acceptable compromise. In any case, one should plan to have as many of these different functions (i.e., packet filters, proxies, and stateful inspection) available as possible, applying each where appropriate.

Ideally, the design of your firewall architecture should precede firewall hardware and software selection. However, we recognize that in some organizations, some form of firewall may already be in place.

Why this is important

Your ability to enforce your organization’s security policies accurately can be severely impaired if you have not chosen an appropriate and effective firewall architecture. This design will determine which policies can and cannot be enforced, as well as how well the firewall will accomplish its objectives over time. Firewall architectures are difficult and expensive to change after deployment, so there is considerable value (cost savings) in creating an effective, scalable, and manageable design first.

Firewall systems provide a policy enforcement mechanism at a security domain boundary. If an adversary can exploit another less protected boundary to gain access into your network (for example, a modem on a user workstation or via a partner’s network), then any firewall systems you have deployed on other boundaries to control access to that network will be ineffective.

How to do it

➤ *Document the environment*

The generation and use of diagrams are extremely important while designing your architecture. They are a good communications mechanism and they are excellent tools to help you avoid mistakes later. The basic rule of thumb is “if you cannot draw it, you cannot build it.” Do not skip or scrimp on this step.

One effective method is to use an electronic whiteboard with a group of knowledgeable people to generate candidate diagrams.

➤ *Select firewall functions*

Firewall functions available in today's products include packet filtering, application proxies, and stateful inspection filtering. Each of these functions implies a certain range of possible choices for deployment platforms. A firewall deployment platform is the combination of the particular hardware and operating system on which the desired firewall functions execute. In some cases, the choice of function and platform can be made independently and in others, the choice of one forces a choice in the other. The following sections describe each of these functions and the platform choices available.

Packet filtering

Since routers are commonly deployed where networks with differing security requirements and policy meet, it makes sense to employ packet filtering on routers to allow only authorized network traffic to the extent possible. The use of packet filtering in those routers can be a cost-effective mechanism to add firewall capability to an existing routing infrastructure. As the name implies, packet filters specify packets to filter (discard) during the routing process. These filtering decisions are usually based on contents of the individual packet headers (e.g., source address, destination address, protocol, port). Some packet filter implementations offer filtering capabilities based on other information, but we consider these under the heading of stateful inspection described below.

Generally speaking, packet filtering routers offer the highest performance firewall mechanism. However, they are harder to configure because they are configured at a lower level, requiring you to have a detailed understanding of protocols².

Packet filtering is typically implemented on two kinds of platforms

- general purpose computers acting as routers
- special purpose routers

The following table shows the principle advantages and disadvantages of each platform.

	General purpose computer acting as a router	Special purpose router
Advantages	Unlimited functional extensibility	Highest performance Large number of interfaces
Disadvantages	Moderate performance Small number of interfaces OS vulnerabilities	Minimal functional extensibility May require more memory

We have found that cost is not a major consideration in choosing a platform for packet filtering.

Special purpose router vendors have added packet filters to their router products to provide limited access controls as a result of customer demand and minimal implementation effort. However, they are router vendors, not security product vendors, so

2. The difficulty primarily arises because of how quickly the rule sets grow in complexity. Low-level debugging is hard and you have to know, in detail, how the router parses the rules you have defined.

when they make a design tradeoff between routing functionality and security functionality, they choose routing. In this context, performance is a routing functionality issue, not a security issue, so it always ranks near the top of the list of design priorities for these routers. In addition, adding filtering to a router

- can negatively impact routing, and therefore networking, performance
- may require additional memory

General purpose computers and the operating system software that runs on them are not typically designed to act as high performance routers, with or without packet filtering. The most common reasons for choosing a general purpose computer include:

- using firewall mechanisms in addition to packet filtering on the same host
- existing in-depth knowledge of the chosen platform
- eliminating filtering load on a special purpose router
- availability of source code

Application proxies

An application proxy is an application program that runs on a firewall system between two networks (see figure 1-3 at the end of this section). The host on which the proxy runs does not need to be acting as a router. When a client program establishes a connection “through” a proxy to a destination service, it first establishes a connection directly to the proxy server program. The client then negotiates with the proxy server to have the proxy establish a connection on behalf of the client between the proxy and the destination service. If successful, there are then two connections in place: one between the client and the proxy server and another between the proxy server and the destination service. Once established, the proxy then receives and forwards traffic bi-directionally between the client and service. The proxy makes all connection-establishment and packet-forwarding decisions; any routing functions that are active on the host system are irrelevant to the proxy.

As with packet filtering, application proxies are available on both special purpose proxy machines and general purpose computers. Generally speaking, application proxies are slower than packet filtering routers. However, application proxies are, in some ways, inherently more secure than packet filtering routers. Packet filtering routers have historically suffered from implementation flaws or oversights in the operating system’s routing implementation on which they depend. Since packet filtering capabilities are “add-ons” to routing, they cannot correct or compensate for certain kinds of routing flaws.

As a result of making more complex filtering and access control decisions, application proxies can require significant computing resources and an expensive host upon which to execute. For example, if a certain firewall technology running on a UNIX platform needs to support 200 concurrent HTTP sessions, the host must be capable of supporting 200 HTTP proxy processes with reasonable performance. Add 100 FTP sessions, 25 SMTP sessions, some LDAP sessions, and some DNS transactions and you have a host that needs to sustain 500 to 1,000 proxy processes. Some proxies are implemented using kernel threads (which can dramatically reduce resource requirements) but resource demands remain high.

Stateful inspection or dynamic packet filtering

We use the terms stateful inspection or dynamic packet filtering to refer to a more capable set of filtering functions on routers. Packet filtering is restricted to making its filtering decisions based only on the header information on each individual packet without considering any prior packets. Stateful inspection filtering allows both complex combinations of payload (message content) and context established by prior packets to influence filtering decisions. As with packet filtering, stateful inspection is implemented as an “add-on” to routing, so the host on which the stateful inspection function is executing must also be acting as a router.

The principle motivation for stateful inspection is a compromise between performance and security. As a routing “add-on,” stateful inspection provides much better performance than proxies. It also provides an increase in the level of firewall function than simple packet filtering. Like proxies, much more complex access control criteria can be specified and like packet filtering, stateful inspection depends on a high quality (i.e., correct) underlying routing implementation.

Refer to “Stateful Inspection Firewall Technology Tech Note.” [Check Point 98] and “Application Gateways and Stateful Inspection: A Brief Note Comparing and Contrasting.” [Avolio 98] for more information about stateful inspection and dynamic packet filtering. Additional information on all firewall functions and the pros and cons of each can be found in *Firewalls and Internet Security* [Cheswick 94], *Building Internet Firewalls* [Chapman 95], *Firewalls Complete* [Goncalves 98], *Third Annual Firewall Industry Guide* [ICSA 98], and *Internet Security Policy: A Technical Guide* [NIST 98]. A recent summary of thirteen vendor firewall products and the functions they support can be found in “Firewalls Market Survey” [SC 99].

We recommend the following as a guideline for choosing firewall functions:

Function	Packet filtering (PF)		Application proxies (AP)		Stateful inspection (SI) and packet filtering		Packet filtering and application proxies		Stateful inspection, packet filtering, and application proxies	
	SP	GP	SP	GP	SP	GP	SP	GP	SP	GP
Protocol / service ^b	A	A	S	S	A	A	A	A	A	A
Support ^c	T	S	T	S	T	S	T	S	T	S
Security requirements ^d	L	L	H	H	M	M	L (PF) H (AP)	L (PF) H (AP)	L (PF) M (SI) H (AP)	L (PF) M (SI) H (AP)
Performance / scale requirements ^e	H	H	L	L	M (SI) H (PF)	M (SI) H (PF)	L (AP) H (PF)	L (AP) H (PF)	L (AP) M (SI) H (PF)	L (AP) M (SI) H (PF)

a. SP - special purpose computer; GP - general purpose computer

b. A - any; S - only specific protocols or services

c. T - turnkey support via vendor; S - site-supported

d. L, M, H - low, medium, high

e. L, M, H - low, medium, high

► *Select the firewall topology*

While the firewall functions described above can be deployed in a wide variety of ways, there are a small number of commonly deployed architectures. They are presented in order of increasing effectiveness.

Basic border firewall (See figure 1-4 at the end of this section.) This is the starting point for all firewalls. A basic border firewall is a single host interconnecting an organization's internal network and some untrusted network, typically the Internet. In this configuration, the single host provides all firewall functions.

Untrustworthy host (See figure 1-5 at the end of this section.) To the basic border firewall, add a host that resides on an untrusted network where the firewall cannot protect it. That host is minimally configured and carefully managed to be as secure as possible. The firewall is configured to require incoming and outgoing traffic to go through the untrustworthy host. The host is referred to as untrustworthy because it cannot be protected by the firewall; therefore, hosts on the trusted networks can place only limited trust in it.

DMZ network (See figure 1-6 at the end of this section.) In a DMZ network, the untrusted host is brought "inside" the firewall, but placed on a network by itself (the firewall host then interconnects three networks). This increases the security, reliability, and availability of the untrusted host, but it does not increase the level of trust that other "inside" hosts can afford it. Other untrustworthy hosts for other purposes (for example, a public web site or ftp server) can easily be placed on the DMZ network, creating a public services network.

Dual firewall (See figure 1-7 at the end of this section.) The organization's internal network is further isolated from the untrustworthy network by adding a second firewall host. By connecting the untrustworthy network to one firewall host, the organization's internal network to the other, and the DMZ between, traffic between the internal network and the Internet must traverse two firewalls and the DMZ.

In each of these architectures, firewalls are used to control access at the border of your network mainly for the purpose of protecting your network from an untrusted network. Firewalls deployed entirely within your network can also be used to provide mutual protection among subnets of your network. Controlling access between internal subnets is no different than controlling access between your network and the Internet, so all of the above architectures can be used as internal firewall architectures as well.

Additional information on these firewall architectures and their pros and cons can be found in *Firewalls and Internet Security* [Cheswick 94], *Building Internet Firewalls* [Chapman 95], *Firewalls Complete* [Goncalves 98], "Firewalls fend off invasions from the Net." [Lodin 98], and *Internet Security Policy: A Technical Guide* [NIST 98].

► *Perform architectural trade-off analysis*

Firewalls are typically thought of in their restrictive or protective sense. That is, they protect your network from the Internet or they restrict access to your network from the Internet. In today's Internet-enabled organizations, firewalls are more frequently thought of as safely empowering the organization to interact with the Internet. As such, firewalls are very much part of an organization's mission-critical infrastructure and they need to be designed accordingly.

As a result, you must make the same architectural tradeoffs in designing your firewall that are commonly made in other mission-critical systems. Architectural characteristics that must be considered include

- performance
- availability
- reliability
- security
- cost
- manageability
- configurability
- function

Refer to “The Architecture Tradeoff Analysis Method.” [Kazman 98] and “Experience with Performing Architecture Tradeoff Analysis.” [Kazman 99] for more information on performing architectural tradeoffs.

Areas to consider include

- availability. Availability is achieved by a combination of reliability and redundancy. Start by choosing hardware and software components that are reliable. If the level of reliability achieved is insufficient, consider using redundant components to meet availability requirements.³
- performance. Based on the anticipated traffic through the firewall system, you may need multiple firewall hosts to distribute the load and handle traffic at an acceptable rate.
- security. Weigh the use of single versus dual firewall systems at your network perimeter. The factors to consider include
 - having outside traffic passing through two firewall systems instead of one (benefits vs. cost)
 - your ability to monitor traffic and the monitoring locations
 - your ability to recover from compromises including disconnecting one firewall system while keeping the other operational
 - your needs for and number of network ports
 - performance
 - failure characteristics
 - expense
 - complexity of firewall system operations and maintenance
 - using multiple firewall systems from different vendors to reduce your exposure to vulnerabilities inherent in a single product (survivability through diversity)

3. A hot standby system provides the capability to automatically and immediately switch workload from the primary system to the standby system. A warm standby usually requires some reconfiguration before the workload can be switched from the primary system. Cold standbys are started from a shutdown state and need extensive configuration upgrades before being used. Having a hot standby firewall system will minimize downtime and maximize flexibility in being able to test broken systems offline.

➤ *Protect your firewall system from unauthorized access.*

If you need to administer your firewall systems remotely, you must use strong authentication and data encryption technologies to prevent adversaries from compromising your firewall systems. The firewall administrator should be authenticated using technologies such as one time passwords or recognized cryptographic protocols rather than using clear text passwords or replayable authenticators. All administrator communications to and from the firewall systems must be strongly encrypted. Consider strongly encrypting any sensitive information (such as passwords, configuration data) stored on the firewall system or on all administrative systems (such as the network management system).

Ensure that you have appropriate physical access controls for the work areas that house the consoles for your infrastructure management and administration systems. Unauthorized users who have physical access to these systems could use them to access your firewall systems. Ensure you have equivalent physical access controls for the work areas that house your firewall system consoles.

Policy considerations

Your organization's networked systems security policy should include

- the risks you intend to manage with the firewall
- the services you intend to offer to untrusted networks from your protected network. These could be offerings to the Internet or to other internal networks.
- the services you intend to request from untrusted networks via your protected network. These could be requests to the Internet or to other internal networks.
- the objective that all incoming and outgoing network traffic must go through the firewall (i.e., that no traffic which bypasses the firewall is permitted, for example, by using modems) — or conversely, that specific loopholes are permitted and under what conditions (e.g., modems, tunnels, connections to ISPs)

In the offering and requesting of services, your policy should ensure that you only allow network traffic

- that is determined to be safe and in your interests
- that minimizes the exposure of information about your protected network's information infrastructure

For additional information on policy-related topics, refer to *Firewalls Complete* [Goncalves 98].

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p053.html>

Figure 1-1: Example of single layer architecture




Figure 1-2: Example of multi-layer architecture




Figure 1-3: Application proxy

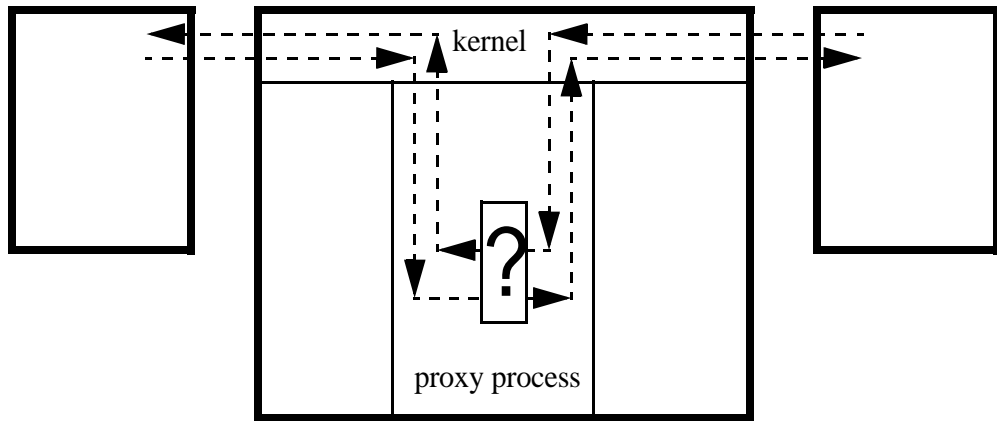


Figure 1-4: Basic border firewall architecture




Figure 1-5: Basic firewall with untrustworthy host architecture




Figure 1-6: Basic firewall with DMZ network architecture





Figure 1-7: Dual firewall with DMZ network architecture



2

Acquire firewall hardware and software.

You need to ensure that you have all of the hardware and software necessary to install, test, operate, monitor, and audit your firewall system prior to its deployment. In addition, you need to ensure that you have adequate physical space to accommodate the equipment and that it can be connected properly in both its test and operational states. You need to seek expert advice if you are unfamiliar with the hardware, any aspects of its configuration, the software, or the physical environment in which it will operate.

Why this is important

You cannot operate your firewall mechanisms effectively, or perhaps at all, if key hardware or software components are missing. If you do not ensure that you have all components on hand prior to deployment, you are likely to experience delays in ordering and acquiring missing components. This could increase the time it takes to deploy your firewall system.

How to do it

➤ *Determine required hardware components*

These may include

- appropriate processors on which to run the firewall software with sufficient processing speed to meet performance requirements
- adequate RAM to meet performance requirements
- devices necessary for software installation (e.g., CD-ROM, floppy drives, keyboard, display, mouse)
- adequate hard disk space to accommodate the operating system, the firewall software, and additional requirements such as log files.
- firewall client administration workstation(s)
- network interface cards
- backup devices and media
- physical space such as rack mount space
- appropriate power (e.g., plug strips, redundant power supplies, continuous power)
- appropriate cabling (e.g., network and console cables)
- testing devices (e.g., network traffic generators and monitors)
- surrounding network infrastructure (e.g., routers, switches, and hubs)
- telecommunications facilities
- spare parts as required

Processor, memory, and disk capacities should be determined on a cost/benefit basis. You want to order the maximum that you can afford. Firewall software processing is typically very resource intensive and you will continue to require increased capacity as your network grows and as your traffic or security needs increase.

Ensure there are sufficient adapter slots for all of the networks that will connect to your firewall system in both test and operational modes. Ensure that they operate at the data rates you require. If you have a very high traffic site, you may need to consider multiple parallel gateways with automatic load balancing so that your firewall systems do not become a bottleneck.

Ensure that you have sufficient spare equipment on hand to meet your firewall redundancy, availability, and failure recovery requirements. For example, if you plan to maintain a hot standby or backup of your firewall system, you need sufficient equipment to operate a fully-redundant system.

➤ *Determine required software components*

These may include

- host operating systems
- patches and fixes to secure the operating system and bring it up to the most current version
- device drivers for all adapters and interfaces required
- any tools that are required to perform software reconfiguration
- firewall software components
- support utilities
- network monitoring tools such as tcpdump to view network traffic during testing and operations
- patches and fixes to secure all software components

➤ *Determine required testing components*

In the same way that the operational environment for the firewall system must be designed, so must the test environment. Refer to “8. Test the firewall system.” for information about determining testing requirements. The test environment should be designed to be as realistic as possible without running the risk of compromising your operational network or the firewall systems under test.

Hardware components may include equipment that serves the role of some or all of the networks that the firewall systems will eventually interconnect as well as equipment used for the purposes of generating simulated traffic.

Software components may include tools to simulate network traffic that will exercise firewall rules.

While it is theoretically possible to exhaustively test a firewall policy by generating and monitoring network traffic, it is practically infeasible. Therefore, a traffic sampling technique must be used. Two possible approaches are to capture or replay existing traffic or generate simulated traffic.

We recommend generating simulated traffic for the following reasons:

- You can exercise traffic of most interest at any point in time by choosing what traffic you generate.
- You are not distracted by traffic that is irrelevant to the test you are currently conducting.
- You do not need to characterize the captured traffic to ensure it adequately covers your areas of interest.
- You do not need to sanitize traffic as it does not represent actual communication.

Most approaches to firewall testing are likely to include a review of log files and the use of network traffic generators and sniffers. Refer to “7. Configure Logging and Alert Mechanisms” for more information about logging practice.

➤ *Acquire all components*

Ensure that you have all hardware and software components available before attempting firewall system deployment.

Conduct a preliminary installation of the firewall software and operating system on the target hardware to ensure that nothing is missing. It is particularly important that you do this upon receipt of hardware and software if your deployment is delayed. If something is missing, you have time to correct the omission before deployment deadlines. If you skip this step, you may not realize the omission until much later. Plan to do this type of preliminary installation many times. The more comfortable you are with the installation process, the more quickly you can perform major reconfigurations or recoveries.

If your firewall operating system (OS) resides in nonvolatile memory (e.g., flash memory), make sure that you can erase its contents completely and rewrite the OS image onto the hardware. Do this for both your primary and all spare OS hardware. This will ensure that your OS hardware works correctly and that you can load a new version of the operating system once the firewall system is deployed.

If you have limited experience with the target hardware or operating system, bring in a knowledgeable consultant or vendor. Document your understanding, the actions that they take, and the recommendations that they make. Have the consultant/vendor sign this document in the event you encounter problems in the future. It may give you some leverage to have them return without incurring additional expense. Refer to “3. Acquire firewall documentation, training, and support.”

Other information

Be aware that installing and configuring firewall hardware and software are difficult and complex processes. Each firewall product is different. It is critically important that you carefully read and understand all documentation provided by product producers. For example, some products expect specific hardware (e.g., graphic adapters) or specific software patches to be present for a successful installation.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p054.html>

3 ***Acquire firewall documentation, training, and support.***

Depending on the firewall architecture you design, you may need some level of training or vendor support when you are deploying a new firewall system. There are a range of choices you need to evaluate in order to determine your requirements for these information sources.

Why this is important

If you are unfamiliar with the technologies that make up your new firewall, you are likely to make potentially costly mistakes. This can cause delays in all aspects of installing, configuring, deploying, operating, and maintaining your firewall system. While the most serious mistake results from incorrect security configurations (exposing your network to a range of possible consequences), even maintaining the underlying hardware and software can be complex enough to warrant training or support.

How to do it

➤ *Determine your training requirements*

You can almost always acquire training services and materials on the various technologies that make up your firewall system from the firewall product vendor. You may also be able to acquire what you need from other organizations that specialize in such training. Start by assessing the skills available within your organization. If staff with the requisite skills are not available, the best way to understand what is needed is to ask the personnel who are candidates for the training; they are most likely to know what they don't know. Assess existing skills and plan to supplement them as necessary in the following areas:

- TCP/IP protocols, services, and routing
- network architecture
- hardware on which the firewall runs or depends
- software on which the firewall runs or depends, including the operating system
- the firewall software
- network security and survivability
- network monitoring
- system management techniques
 - installation
 - maintenance
 - backup and recovery
 - system security
 - auditing, logging, and monitoring

Training is a relatively expensive commodity; make sure you get what you need — no more, no less.

Be sure to consider a range of training delivery methods such as

- classroom including on-site instruction and the use of web-based or video-conferencing technologies
- self-paced (conventional or computer based)
- books and manuals
- journals and magazines
- conferences and user groups
- World Wide Web resources

Make sure to consider your future training requirements, including those for new personnel, in your plan.

Depending on the extent of your requirements as described above, we recommend that you schedule training well in advance of any firewall deployment activities but close enough to the start of deployment to be applied.

➤ *Determine your support requirements*

Vendor support may be essential when you are trouble shooting complex problems. Vendor support can also be used in lieu of training to address specific questions whose answers are not clear or present in available documentation.

Vendor support is generally negotiated in the form of a service level agreement¹. It can come in several forms: access-controlled Web site support, phone support, and onsite support. Phone support will likely have the following characteristics

- a choice of 24 hours a day, seven days a week or during your normal business hours (Monday - Friday, 8:00 a.m. to 5:00 p.m.)
- a number of specified individuals who can call for support or an unlimited number of callers
- a specified calendar time period for the support (one month, six months, one year)
- the use of a more in-depth service or onsite support if those that perform vendor phone support cannot solve your problem
- a flat monthly or annual rate for phone support; an hourly rate for in-depth or onsite support that is billed as you use the support

As with training, be sure to consider product support in addition to the firewall system such as support for the operating system.

We recommend that you obtain vendor support as soon as you have selected your firewall systems and before you actually start testing and deployment.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p055.html>

1. For more information on developing a service level agreement, refer to <http://www.gtlaw.com.au/pubs/negotiating.html> and <http://www.gtlaw.com.au/pubs/negotiating-service.html>.

4 *Install firewall hardware and software.*

Install and configure the operating system that will execute the firewall software followed by installing and configuring the firewall software. These two steps should be performed on the firewall hardware you intend to use in your production environment but deployed in the test environment and configuration (Refer to “8. Test the Firewall System” for information on using a test configuration). You need to ensure that all hardware and software are properly configured and operate as expected to the extent possible in the test configuration.

You need to configure the operating system on your firewall host in the minimum essential configuration so that only those services necessary for firewall operation and maintenance are included. You need to include all applicable patches or fixes for both the operating system and the firewall software.

Why this is important

The most common cause of firewall security breaches is misconfiguration of the firewall system. Various references on penetration testing show that well over half of the firewall systems regularly tested are not properly configured. According to ICSA¹, seventy percent of sites with certified commercial firewalls are still vulnerable to attacks due to misconfiguration or improper deployment.

Exercising your installation and configuration procedures in a test environment will allow you to learn the requirements to efficiently install and configure both the operating system and your firewall software while minimizing the impact on your operational systems. It will highlight what, if any, hardware may be missing in your initial configuration.

If you do not install the operating system and your firewall software with a minimal service configuration and with all applicable patches, you risk

- exposing your organization’s network to intrusions that exploit well-known vulnerabilities for which patches exist
- not being able to get support from your vendor. Vendors almost always require the underlying system to be current before they will answer questions.
- not having a stable platform on which to run the firewall software. Many patches are related to reliability and recovery.

1. January, 1999 quote.

How to do it**➤ *Install a minimum acceptable operating system environment***

Ensure that your firewall system configuration includes only those packages and services that are required for firewall system operation.

This can be accomplished by either

- removing all software that is not needed (if this can be determined) after installation
- including only that software which is needed, selectively adding specific packages and services back in as you determine that they are required

Examples of services that are typically included in a default operating system configuration that should be removed are X Windows services, telnet (assuming ssh is installed and configured), NFS for Unix-based operating systems, and NetBios for Microsoft NT operating systems.

For some firewall products, the process of installing the firewall software will force a minimal configuration of the operating system such as removing unnecessary services if you did not do this before the installation process.

Keep in mind that packet filtering functions typically run in the operating system kernel (for performance reasons) and, therefore, packet filtering software is fairly sensitive to a specific kernel version and release number.

Once you are satisfied that the operating system and the firewall software is successfully installed, you should repeat the sequence to ensure that the process can be done again. The second time, document it. The third time, have an outside person who was not involved in the first two installations follow the documentation to see if it is correct and complete.

Take appropriate steps to ensure that any redundant systems are in a state consistent with the systems to be used in production. Ensure that you can easily switch between your primary firewall system and any redundant systems.

Your installed environment may not have all of the necessary troubleshooting and support tools necessary to determine what has happened if anything goes wrong during the installation process. You may need to install the firewall system on another host that has better diagnostic tools if you run into problems. After you understand the problems and know how to compensate for them, you can complete the installation on the production hardware.

➤ *Install all applicable patches*

This information is available from your operating system and firewall software vendors. Determine how to deliver patches securely to the firewall system. Some products require that you do this using removable media (disk, CD-ROM), not via a network.

As an operational consideration, if your redundancy requirements result in your having an identical hot backup or standby firewall system (which we strongly recommend), you can consider installing and testing any new patches on the redundant system and then switch from the current operational system to the redundant system.

Your vendor service level agreement should state that the firewall software will always be fully functional if all of the operating system patches are installed.

In addition, you need to ensure that those responsible for your firewall operating system and software have time set aside to periodically review applicable public and vendor information sources for security patch updates. These sources regularly report current intruder trends, new attack scenarios, security vulnerabilities, methods for their detection, and guidance to address them.²

➤ *Restrict user and host access*

The only users who should have access to your firewall system are the firewall system administrator, those authorized by policy, and individuals involved in operating and maintaining your information technology infrastructure.

For some firewall products, the process of installing the firewall software will automatically disable access to the firewall system by all users (except those mentioned above) if you have not already disabled their access before installation.

We recommend that you allow remote access to your firewall system only via mechanisms that are strongly authenticated and strongly encrypted, even on your organization's internal networks. Some firewall products provide the capability to restrict the administrative client to a specific IP address and a specific port. We do not believe that this is adequate security; encryption is required as well. IP addresses and ports are too easily spoofed.

➤ *Disable IP forwarding*

Make sure packet forwarding is disabled until after the firewall software is operational.

While booting firewall hosts, there may be an interval of time after the operating system is functional, including networking, but the firewall software is not yet functional. During this interval, packets may flow freely through the firewall system. Make sure that no packets are forwarded before the firewall software is functioning by doing one or more of the following

- disable IP routing before any interfaces are enabled
- do not enable network interfaces before the firewall software is functional

➤ *Backup your system*

When installation is complete, perform a backup of the entire firewall system. Use this backup to restore the production system (or one identical to the test firewall system) for operation. Verify that both the operating system and the firewall software operate properly from the restored backup version. Refer to *Securing Desktop Workstations* [Simmel 99], specifically the practice “Configure computers for file backups.”

2. A list of information sources can be found at the end of the Executive Summary of this module and at <http://www.cert.org/security-improvement/implementations/i040.01.html>, titled “Maintaining currency by periodically reviewing public and vendor information sources.”

Policy considerations

Your organization's networked systems security policy should require

- timely evaluation, selection, and installation of patches and other corrections that you need to operate securely
- that only authorized personnel have access to the firewall system via authorized, strongly authenticated mechanisms
- that firewall system installation is performed in an environment isolated from your operational networks
- that your firewall system is backed up on a regular basis

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p056.html>

5 *Configure IP routing.*

Routing is the process of deciding the disposition of each packet that a router handles. This applies to incoming packets, outbound packets leaving your network for external destinations, and those packets being routed among your internal networks.

There can be only two dispositions: forward or discard. The routing mechanism decides between these two using the destination IP address in the packet header. This decision process is governed by a data structure called the routing table. The routing mechanism should not be used to implement security policy; it is too dynamic and unreliable. Routing functions and supporting structures are designed to route packets efficiently and reliably, not securely.

A firewall system's routing configuration reflects its view of the topological configuration of the networks to which it is attached. Most firewall systems' routing configurations rarely change; they are static¹.

Why this is important

You should have a routing configuration that reflects your network topology so that your firewall system will be able to deliver legitimate packets to their desired destinations.

How to do it

➤ *Obtain IP addresses*

Obtain a unique IP address for each interface on each firewall system.

Each network to which a firewall system is attached has a procedure to obtain new IP addresses. For the Internet, this is obtained from the Internet Service Provider (ISP) that will connect to your firewall. For internal networks, including any DMZ networks you intend to establish, you must obtain IP addresses from within your organization.

➤ *Establish routing configuration*

A firewall system's routing table contains a list of IP network addresses for which the firewall system is intended to provide routing services. Each row of the table describes one network. The index used to access a row in the table is the destination network address of the packet currently being routed. If table lookup is successful, the table provides either the address of the next router to which to send the packet or the interface to use to send the packet out. That next router is used as the intermediate destination and the packet is forwarded there. If the table lookup fails, the packet is discarded. An ICMP "unreachable"

1. Dynamic updates to routing configurations do occur. However, we state them as static here because the large majority of firewalls today have two interfaces — one to the Internet and one to the organization's internal network. In this case, routing is static. Most ISPs handle all dynamic routing, presenting a static interface to their customers' systems at all times.

message may be returned to the source indicating that the packet was undeliverable.

If you are replacing an existing firewall system or router, thoroughly examine the routing configuration of your system to determine the network topology that it describes. Ensure, as a first step, that the routing configuration of the new firewall system is consistent with your current system before departing from that configuration.

Information on routing protocols and the process of establishing a routing configuration is fully described in *Internetworking with TCP/IP, volume 1: principles, protocol and architecture. 3rd Edition* [Comer 95].

Policy Considerations

Your organization's networked systems security policy should

- require that configuring IP routing for your firewall system is performed in an environment isolated from your operational networks
- specify what connectivity is to be permitted with the specific statement that all other connectivity is denied

Other information

Your routing configuration is derived from your network topology. You should not attempt to implement aspects of your security policy with routing. It is too imprecise, exercises insufficient control of your incoming and outgoing network traffic, and has no support for auditing and logging.

If your firewall design is based on a multiple layer firewall architecture with a DMZ so that all inbound and outbound packets travel through both firewall systems, you need to take the following into account:

- For the outside firewall system (the one that sits between the external world and the inside firewall system), the routing configuration is more complex and the packet filtering rules are more simple. Formulating the routing configuration and the filtering rules needs to be done separately and somewhat in sequence.
- For the inside firewall (the one that sits between the outside firewall and your internal networks) the routing configuration is more simple and the packet filtering rules are generally more complex (depending on the policy you are implementing). You need to formulate the routing configuration and the filtering rules concurrently.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p057.html>

6 *Configure firewall packet filtering.*

Packet filtering is the process of deciding the disposition of each packet that can possibly pass through a router with packet filtering (firewall host)¹. For this discussion, assume there are only two dispositions: accept and reject. IP filtering provides the basic protection mechanism for a routing firewall host, allowing you to determine what traffic passes through it based on the contents of the packet, thereby potentially limiting access to each of the networks controlled by the firewall.

The criteria used in each filtering rule for determining the disposition can be arbitrarily complex. For a router with packet filtering, there may be multiple points in the routing process where the rules are applied; typically, for arriving packets, they are applied at the time a packet is received and, for departing packets, they are applied immediately before a packet is transmitted. There may be different rule sets at each point where filtering is applied.

Your firewall's packet filtering rules should implement some portion of your organization's network security policy. If the entire policy can be implemented in packet filters, then other firewall mechanisms may not be required. If some elements of your policy can not be implemented with packet filters, then consider additional firewall mechanisms such as proxies.

Why this is important

Packet filtering can be a high performance, low cost means by which to implement a substantive portion of your network security policy.

How to do it

➤ *Design the packet filtering rules*

The criteria that are available to influence the content of the rules you design comes from two sources: information intrinsic to a packet and information extrinsic to a packet. Intrinsic information is information that is contained within a packet itself. Examples of intrinsic information include

- packet header information
 - source address
 - destination address
 - protocol
 - source port
 - destination port
 - packet length

1. Filtering uses a range of information in the packet header (for example, source and destination IP addresses, port, and protocol) while routing uses only the destination IP address.

- connection state information
- packet payload (message content)

Some or all of this information may be referenced in filter rules depending on the actual implementation of the packet filter mechanism in the firewall software you have selected. Most packet filters only support references to information in headers.

Extrinsic information is information that exists outside of a packet. Examples of extrinsic information are

- arrival interface on the router or firewall
- departure interface on the router or firewall
- context maintained by the firewall software that pertains to a packet (for use in stateful inspection)
- date and time of packet arrival or departure

Packet filters cannot generally reference extrinsic information. Most packet filters are implemented as separate sets of rules for each interface, sometimes with separate sets for arriving and departing packets. By placing a given rule in the appropriate interface's rule set, you are using extrinsic information in the design of your rules.

When designing rules for a new network interconnect (one that does not currently exist in firewalled or unfirewalled form), we recommend the following approach:

1. You can generally assume that the last rule in every rule set of every firewall system is to deny all packets. However, we recommend that you explicitly add this rule to remind you that this is the policy you are implementing and to express the rule set more completely.
2. Design anti-spoofing rules and put them at the top of each rule set².
3. Canvass the potential users of this interconnect to find out what they expect to be doing. Collect this list into a table of protocols, ports, and source and destination addresses. Select those that implement your security policy.
4. Sort the table by protocol and then by port.
5. Collapse the matching protocols rows and the consecutive ports rows together into one new row that specifies a range for the port.
6. Convert this table into a set of rules and insert them between the anti-spoofing rules and the "deny all" rule at the end of the rule set.

You need to take the following additional guidelines into account:

- For some firewall software, installing an empty set of packet filtering rules does not mean that all traffic is allowed to pass through. It instead instructs the firewall to block all traffic (including firewall configuration session packets). If this happens, you will need physical access to the firewall management console to recover the system.
- Watch for default packet filtering rules that may come with the firewall software. These are sometimes not obvious and will likely not be in compliance with your security policy. Carefully review your firewall software documentation.

-
2. Refer to <ftp://ftp.isi.edu/in-notes/rfc2267.txt> for more information about spoofing and designing anti-spoofing rules.

- If your firewall has separate rule sets for arrival and departure on each interface, repeat the rules in the arrival rule set of each interface and in the departure interface of the others. This reduces the possibility of an oversight.
- If your firewall only has two interfaces, then you can automatically generate rules for one interface from the other. This becomes more difficult as the number of interfaces increases, eventually becoming computationally infeasible.
- For anti-spoofing rules to work as intended, it is necessary for your firewall to be able to distinguish between arrival and departure on each interface independently. You need to specify sets of rules that reference interfaces and direction; otherwise, you cannot implement anti-spoofing rules without interfering with other rules.
- Check to see if your firewall system has the capability to generate anti-spoofing rules automatically from routing tables. The advantage of this approach over manually designed rules is that the rules may be able to adapt automatically to routing changes. This is important if your firewall has multiple routes to the same destinations (e.g., multiple connections to the Internet). If you do not automatically generate anti-spoofing rules, then routing changes may adversely affect your ability to reach certain destinations since routing and filtering could then be in conflict.³
- Remember that packet filters must be based on IP addresses, not host names or use of DNS services. If your firewall software allows you to specify addresses in the form of host names, the names must be converted to addresses and these addresses are then used in filters before filters can take effect. If a host's address changes, then the firewall will not recognize the change until the filter configuration is reloaded and the names are converted to addresses again.
- If your filtering software allows you to refer to established TCP connections, you may wish to use this feature to allow return traffic to be forwarded for established connections. This is a useful feature to allow packets coming back from TCP connections originating within your internal network to the Internet into your network. Note that this is not foolproof. It is possible for an adversary to forge such packets, thereby getting them through your firewall and into your network. If all of your internal systems have properly operating TCP implementations, they will discard these forged packets. But this is a risk you must decide to accept or reject.
- Given that UDP is a connectionless protocol, there is no concept of an "established" state. Therefore, there is no completely safe way to allow "return" UDP packets — they simply cannot be positively associated with an outgoing packet. However, in spite of its inherently less safe nature, a number of critical services depend on UDP packets. For example, the domain name service (UDP port 53) must be functional in order for your internal hosts to operate correctly. For each such UDP service, it may be possible for you to determine the level of threat in permitting its packets and thus make a risk management choice of permitting or denying the service. If you decide that you must support a given service, but you cannot accept the risk of permitting that service's traffic through your packet filtering firewall, it may be necessary to consider a higher function firewall mechanism, such as a proxy, stateful inspection, or running separate services on each side of the firewall system (such as having two DNS servers).
- ICMP is also a connectionless protocol. Therefore, it suffers from the same types of filtering limitations and threats as UDP. However, there are only thirteen types of ICMP packets. For each type, you should explicitly decide what you are willing to permit.

3. Be aware that routing topology changes may create conflicts between routing and filtering policies for non-spoofing related filter rules. This is the primary reason that most firewalls operate with a static routing configuration.

- Source routing is a function of IP routing that allows the packet originator to influence routing decisions as the packet traverses networks. We recommend that you disable all source routing functions in your firewall's router and that, if possible, you deny any packets that have specified source routing options.

Refer to *TCP/IP Illustrated, Volume 1: The Protocols* [Stevens 94] for details on TCP, UDP, and ICMP.

➤ *Document the packet filtering rules*

Document packet filtering rules so that anyone reading the documentation and configuration data is able to see how each filtering rule enforces its part of the firewall policy. Describe why each rule is included.

Specifying the detailed rules that implement your firewall policy can be a very complex and time-consuming process. It is very common that several rules are required to achieve one integral function. Group the rules that go together and include comments or other documentation that explains what the group does (to the extent that the filter syntax permits).

Most commercial firewall products provide complex configuration managers and user interfaces for specifying rules. Some products provide the capability to aggregate related sets of rules into groups. However, the language used to express the rules is typically network-based, not policy-based. This means that you can specify a rule and understand what it does from a network perspective, but you may still have little idea of the policy implications.

In many current products, firewall system performance degrades as the number of filtering rules increases. So you may be forced to combine rules to optimize performance (make sure to include good documentation to explain what the optimized rule set does and how it performs). This further contributes to obfuscating the rules from a policy perspective. Establishing and maintaining your understanding of the rules from both a policy viewpoint and a network viewpoint at the same time is very difficult.

One strategy to address this complexity is to maintain your firewall policy and the rules that implement it in a policy language of your own design. We know of no products to do this. Developing a language that works for the firewall and system administrators in your organization might be worth the effort as it simplifies the operation and maintenance of your firewall system and reduces the probability of introducing errors over time.

It is easy to make mistakes in altering rules as requirements change. We recommend keeping the rule sets under revision control (using a tool such as RCS⁴ or SCCS⁵) so that you can look back to prior configurations when necessary (and it will be necessary).

➤ *Install packet filtering rules*

Now that you have one or more sets of rules, install them in your firewall test environment. Most firewall software has a mode of operation that allows the installed filters to be dumped to a file for examination. Become familiar with this feature. Install your filters, dump them, and compare the two. Sometimes you will find ambiguities in the input language that result in your implementing something different from what you intended. Continue this *install-dump-compare* cycle until you are satisfied that you have installed what you intended to install.

4. Revision Control System

5. Software Configuration Control System

In actual use, all three steps of this practice (design, document, install) are performed in parallel using a process of successive refinement. It is also common to perform initial rule-specific testing during this activity to ensure that each rule does what you expect it to do. Refer to “8. Test the Firewall System” for more information on testing.

Policy considerations

It is rarely the case that packet filtering can implement an organization’s security policy exactly. Your managers must decide what level of accuracy and precision is required in implementing your security policy. It will almost certainly be the case that you will have to implement rules that are either more permissive or more restrictive than your policy. You need to determine how to handle this discrepancy.

Your organization's networked systems security policy should state

- that all network traffic that is not explicitly permitted should, by default, be denied
- that configuring packet filtering for your firewall system is performed in an environment isolated from your operational networks

Other information

1. Two objectives common to all firewall systems are that you want to only allow network traffic that you have determined is consistent with your security policy (or conversely, disallow traffic that is inconsistent with your policy) and you want to minimize the amount and usefulness of information about your information infrastructure that is disclosed to those outside of the firewall.

These objectives can be achieved in a number of ways. When selecting filtering mechanisms, we recommend the following guidelines:

- Controlling what packets are forwarded can be achieved in two ways. You can deny packets you know are unacceptable and permit everything else or you can permit packets you know are acceptable and deny everything else. Logically, these are equivalent. Operationally, they are quite different. Given that there are a potential of 131,070 TCP and UDP services on any given host, it is operationally infeasible to determine which of these might be acceptable or unacceptable. Fortunately, your users will tell you what is acceptable (usually in the form of complaints). We believe that taking advantage of your users’ suggestions leads to the best practice for packet filtering today, which is to explicitly permit packets which you know to be acceptable (and reflective of your policy) and to deny everything else.
- Your policy should carefully limit the use of your organization’s internal networks and services by users who are not members of your organization. If you are providing services specifically for use by nonmembers (such as access to your public Web site), your policy should require that they be isolated from your internal systems. This decreases your risk even if one of the services is vulnerable to attack because access is strictly controlled.
- Your policy should address access requirements by members of your organization who are located on untrusted networks (e.g., mobile users on the Internet, employees located at a business partner site). You may need to implement mechanisms to permit appropriate access through the firewall system to your internal networks or systems by these people using strong authentication and encryption.

2. There are references⁶ dealing with the dangers of broken IP implementations, ranging from fragmentation reassembly errors to sequence number prediction to incorrect

6. Refer to the list of general firewall references in the Executive Summary of this module.

interpretation of invalid packet headers. Any such IP implementation flaw can compromise the effectiveness of a packet filtering firewall.

3. If your firewall design is based on a multiple layer firewall architecture with a DMZ in a way that all inbound and outbound packets travel through both firewall systems, you need to take the following into account:
 - For the outside firewall system (i.e., the one that sits between the external world and the inside firewall system), the routing configuration is more complex and the packet filtering rules are more simple. Formulating the routing configuration and the filtering rules needs to be done separately and somewhat in sequence.
 - For the inside firewall (i.e., the one that sits between the outside firewall and your internal networks) the routing configuration is more simple and the packet filtering rules are generally more complex (depending on the policy you are implementing). Formulating the routing configuration and the filtering rules needs to be done concurrently.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p058.html>

7

Configure firewall logging and alert mechanisms.

You want your firewall systems to log activities pertinent to firewall operation and the rules the firewall will enforce. For significant firewall events, you want your firewall system to alert you in real time that these events have occurred.

There are two types of logging you need to specify for your firewall systems

- logging associated with the packets arriving and departing the firewall (e.g., packet denied, packet forwarded)
- logging associated with the operation of the firewall software and the system on which it runs (e.g., no more disk space, no more memory, firewall logs are full)

For most firewall software, there are logging options associated with each individual packet filtering rule. Every packet processed by such a rule has these options applied to it for the purposes of logging. An example is to log packets denied by the firewall software. Logging options can also be used to specify the level of logging, which commonly includes real-time alert mechanisms such as paging, electronic mail, or executing an arbitrary program.

Why this is important

The most important reason for logging is to ensure the continued operation of the firewall system. Logged events related to the operational status of the firewall are critically important in preventing and recovering from failures. They are also an important auditing tool to ensure that the proper security configurations (e.g., packet filters, proxies) are installed on the firewall system. Logs of this sort are generally small and can have long term value for a variety of purposes.

Completely apart from firewall operational concerns are concerns about security. Logs can be useful for you to determine how an intrusion might have occurred for the purpose of improving the quality of the firewall implementation. Logs for this purpose have value only over the time period where intrusions can be reasonably resolved — in our experience no more than three to four weeks.

Logs might also be used in intrusion detection. Intrusion detection is the process of detecting attempted, failed, and successful attacks against your network. Intrusion detection is beyond the scope of this practice. Logs for this purpose may have historical value. Refer to *Intrusion Detection: Network Security Beyond the Firewall* [Escamilla 98], *Preparing to Detect Signs of Intrusion* [Kochmar 98], and *Detecting Signs of Intrusion* [Firth 97] for more information on using logs for the purpose of intrusion detection.

Since log files are typically voluminous and difficult for humans to process, you should use alert mechanisms to notify you of any significant event. It is generally impractical to depend on manual analysis of logs to detect significant events.

How to do it

➤ *Design the logging environment*

You need to determine

- the location of firewall log files
 - on the firewall system itself
 - on a remote logging host accessed via a network
- the expected size of log files
- the rate at which data is logged to the log files
- who needs access to the log files and the level of access
- whether or not logging is to be encrypted
- how log files are to be backed up and recovered

We recommend the following as general heuristics to guide your design:

- If you log to a host other than the firewall host, make sure your packet filter logging is configured not to create log entries for log packets. This will create an infinite logging loop. See the next step.
- Log to both local log files on the firewall host and to a central log host. Log as much information as you can locally and keep it for as long as disk space permits. Log only information that you need or will use to the central log host.
- Use a separate logging-only network to log to the central logging host. This prevents logging from affecting normal traffic. You can use this logging-only network to log intrusion detection traffic as well.
- Do not allow disk space requirements to be a constraining factor for logging. Disks are inexpensive so make sure you have enough disk space on both the firewall host and the central log host for local logging.

➤ *Select logging options for packet filter rules*

Design the logging options on each packet filter rule. You need to decide what purpose is served by logging packets affected by the rule. Generally logging for its own sake is not a wise use of resources.

Logging options for each rule may include

- packets that are denied upon arrival at the firewall system
- packets that are denied upon departure from the firewall system
- packets that arrive or depart within a specified time or date interval

In addition to logging options for each rule, you might be able to configure other aspects of logging that relate to packets being filtered. For example, your firewall software may support summarization of individual filter logs. These can be useful for seeing trends and, in general, are more valuable to keep for an extended period than the individual log entries.

Although it is not typical practice to log packets that are permitted by filter rules, it is occasionally interesting to track a particular kind of traffic. For this purpose, you might consider adding a redundant filter rule (one that either permits or denies packets that are already permitted or denied by another rule) just so that you can specify different logging options for that rule. If you do this, make sure to document that fact that the rule is redundant and is not essential to the implementation of your security policy.

If you are replacing an existing firewall system, thoroughly examine the logging options specified by your existing system. Consider retaining those logging options for rules in the new configuration that are consistent with your security policy.

➤ *Design the alert mechanism configuration*

You should design the real-time alert function of your firewall system to ensure that important event notifications are delivered to the appropriate people as quickly as possible with a minimum number of false notifications.

These may include notification of

- unsuccessful user and host login attempts
- packet filters being modified or disabled in the firewall system
- successful logins to the firewall system
- changes to certain files on the firewall system
- operational events (e.g., logs full, memory or disk shortages, system reboots)

If you followed our recommendation above to establish a central log host, then we further recommend that you establish a central alert mechanism that also operates on the central log host. It may be the case that if you choose this approach, the firewall alert mechanism can be disabled and all alerts can be handled by one system.

If you are replacing an existing firewall system, thoroughly examine the alert options specified by your existing system. Consider retaining those parts of the alert mechanism configuration that remain consistent with your current security policy and your firewall system operation policy.

➤ *Acquire or develop supporting tools*

Acquire or develop tools that monitor your log files and summarize the content of your log files beyond those provided by the vendor. These will provide you with the capability to review only those events of interest.

You will need to develop log file archive mechanisms unless your design includes enough disk capacity to retain log records until they are discarded. We recommend that you use WORM (Write Once, Read Many) devices for all archiving.

Policy considerations

Your organization's networked systems security policy should

- require your firewall system to record all significant activity (such as administrative changes and attempts to breach filter rules) by doing thorough auditing and logging
- specify that configuring logging and alert mechanisms for your firewall system is performed in an environment isolated from your operational networks
- require notification to designated administrators of suspicious behavior that can be detected by the firewall system. This would include events related to the firewall system such as failed login attempts and requests to disable filter rules.
- specify guidelines for
 - handling archived log information
 - how long information should be retained
 - discarding log information

Other information

For a fuller treatment of the subject of identifying and enabling logging mechanisms and protecting your log files, refer to *Preparing to Detect Signs of Intrusion* [Kochmar 98], specifically the practice “Identify and enable system and network logging mechanisms.”

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p059.html>

8 *Test the firewall system.*

The purpose of the test activity is to verify that the firewall system works as intended. You should

- plan testing activities to demonstrate that routing, packet filtering, and logging and alert capabilities perform as designed
- test recovery plans for firewall system failures
- design your initial regression testing suite

The features that must be tested include

- hardware (processor, disk, memory, network interfaces, etc.)
- operating system software (booting, console access, etc.)
- firewall software
- network interconnection equipment (cables, switches, hubs, etc.)
- firewall configuration software
 - routing rules
 - packet filtering rules and associated logging and alert options

Why this is important

Testing your firewall system and verifying that it operates properly increase your confidence that it will perform as designed. You should understand the types of failures that are possible for each system component and recovery techniques for each type of failure. This will allow you to exercise your response and recovery processes when and if these failures occur once the firewall system becomes part of your operational infrastructure.

The most common cause of firewall security breaches is a misconfiguration of your firewall system. Knowing this, you need to make thorough configuration testing (of the firewall system itself as well as all of the routing, packet filtering, and logging capabilities) one of your primary objectives.

How to do it

- *Create a test plan*

You need to plan to test both the implementation of the firewall system and the policy being implemented by the system. To test the implementation of the system:

1. Create a list of all field-replaceable system components whose failure would significantly interfere with the ability of the firewall to meet its mission goals.

2. For each such component, create a short list of the most likely modes of failure that would affect firewall operations. De-emphasize components whose failure modes are unlikely to occur or those whose failure modes are likely but their impact on firewall operations is unlikely.
3. For each relevant failure mode
 - Design a test scenario that would either directly cause or simulate the failure.
 - Design a mitigation strategy that minimizes the impact of the failure.

An example of a test scenario is to assume that the host system on which the firewall software runs has an unrecoverable hardware failure which prevents it from performing any packet forwarding, such as might occur if the network adapter failed. A way to test this failure might be to simply unplug the network from the interface to simulate the failure.

An example of a recovery strategy would be to maintain a totally redundant firewall system, switching to the redundant host when failures occur to minimize the time period during which packet transmittal is affected.

Testing the policy installed in the system is more difficult. It is not feasible to exhaustively test an IP packet filter configuration; there are too many possibilities. We recommend that instead of exhaustive tests, you use boundary tests. In these tests, you identify boundaries in your packet filter rules and then sample test the regions immediately adjacent to each boundary. To do this

- for each rule, you identify every boundary in the rule. In general, each constrained parameter in a rule contributes either one or two boundaries. The space being partitioned is a multidimensional packet attribute space. Common attributes include: protocols, source addresses, destination addresses, source ports, and destination ports. Basically, every attribute of a packet that can be independently checked in a packet filter rule defines one dimension of this space. For example, a rule that permits TCP packets from any host to your Web server host on port 80 has checked three attributes (protocol, destination address, and destination port) which partitions the attribute space into three regions: TCP packets to Web server at ports less than 80, port 80, and ports greater than 80.
- for each region, you generate some test traffic which you have engineered to stay within that region. You verify that the firewall either rejects or forwards all traffic for a given region. Within a single region, all traffic should be rejected or forwarded; that is the purpose for partitioning the packet attribute space.

For a complex set of rules, this can be a tedious process and may not be practical. If it is not practical to test the rules, request that several people review them and ask one person to explain to the others what each rule does.

The test plan needs to include the test cases, configurations, and expected results for

- testing the routing configuration, packet filtering rules (including service-specific testing), and logging and alert options
- testing the firewall system as a whole (such as hardware/software failure recovery, sufficient log file sizing, proper archival of logs, performance monitoring)
- exercising both normal conditions and excursion (anomaly) conditions

You also need to describe the tools you intend to use (such as scanners, monitoring tools, and vulnerability detection tools) and the tests you intend to run using those tools.

➤ *Acquire testing tools*

Develop firewall test tools if your firewall product does not come with these capabilities.

Types of firewall test tools include¹

- network traffic generators (such as SPAK (Send PAcKets), ipsend, or Ballista)
- network monitors (such as tcpdump and Network Monitor)
- portscanners (such as strobe and nmap)
- vulnerability detection tools (there are a range of commercial tools available from various vendors)
- intrusion detection systems such as NFR² (Network Flight Recorder) and Shadow³

Refer to *Preparing to Detect Signs of Intrusion* [Kochmar 98], specifically the practice "Identify and install tools that aid in detecting signs of intrusion" and the supporting implementation "Identify tools that aid in detecting signs of intrusion."

➤ *Test the firewall functions in your test environment.*

Establish a test configuration so that your firewall system is interconnected between two isolated hosts, one playing the role of the external world and the other playing the role of your internal hosts. See figure 8-1 at the end of this section. Ensure that the default gateway for the internal host is set to the firewall system under test. If you have chosen an architecture that supports centralized logging (which we recommend), place both the internal host and a log host on your internal network so that you can test logging options. If logging is performed on the firewall host, you can connect the internal host directly to the firewall host.

Have scanning or network sniffing tools in place on your outside and inside hosts to capture all traffic in both directions (inside to outside, outside to inside).

Perform the following steps:

- Disable packet filtering.
- Inject packets that will exercise all routing rules and send these through the firewall system.
- Ensure that packets are routed correctly by examining the firewall logs and your scanner results.
- Turn on packet filtering.
- Inject network traffic that is an appropriate sampling of all possible source and destination IP addresses, across all ports, and for all protocols.
- Ensure that packets intended to be blocked (denied) are blocked. For example, if all UDP packets are to be blocked, ensure that none get through. Ensure that packets intended to enter or exit (permitted) do enter and exit. Do this by examining your firewall logs and scanner results.

1. Information on some of these tools can be found at <http://www.cert.org/security-improvement/i042.07.html> (titled "Identifying tools that aid in detecting signs of intrusion") and at the COAST Web site <http://www.cs.purdue.edu/coast>.

2. NFR can be found at <http://www.nfr.com>.

3. Shadow is an intrusion detection system developed by the U.S. Navy. Information on Shadow can be found at <http://www.sans.org>.

- Scan for open and blocked ports to ensure your firewall system is performing as intended.
- Examine all of the network traffic that is logged and verify that the logging options associated with each packet filtering rule are operating as intended.
- Examine all of the network traffic that is logged and verify that the alert options associated with each logging option are sending alerts to the designated destination (such as the firewall administrator) using the specified mechanism (such as paging or email).

Plan to conduct this step and the next step with at least two people: the original implementer of the routing configuration, packet filtering rules, logging options, and alert options, and an independent person who reviews what has been implemented, understands the intent, and agrees that the network topology and security policy have been reflected correctly.

► *Test the firewall functions in your production environment*

This step assumes that you are migrating from a single layer firewall architecture (see figure 8-2 at the end of this section) to a multiple layer architecture (see figure 8-3 at the end of this section). This step also assumes that you have a network topology of one or more private networks and one or more public networks. The public networks typically connect hosts that respond to internal and external requests for service such as WWW (HTTP), FTP, email (SMTP), and DNS. These hosts may also respond to internal requests for service such as SNMP, file access, and logging. The public network as described here can serve as your DMZ. The private network typically connects hosts that service your internal users including individual user workstations. See figure 8-4 at the end of this section.

Perform the following steps:

- Connect your firewall system to your public and private networks.
- Set the routing configuration on selected public and private network hosts to direct traffic through the firewall system. The basis for selection is on a service-by-service basis, for example, the Web server on your public network and the host storing the files that the Web server needs to access on your private network. Cycle through the selection and exercise of all services such as Web, file access, DNS, mail, and logging. See figure 8-4 at the end of this section.
- Log the firewall system's incoming and outgoing network traffic. Use a scanner or network sniffer to observe what is happening.
- Ensure that packets intended to be blocked (denied) are blocked. For example, if all UDP packets are to be blocked, ensure that none get through. Ensure that packets intended to enter or exit (permitted) do enter and exit. Do this by examining your firewall logs and scanner results.
- Scan all hosts in a selected portion of your network that includes the firewall system. Verify that you cannot gain any undesired information due to the scanning packets being blocked (denied). Attempt source port scanning using a well-known port such as the FTP-data port (port 20) to ensure that you cannot use the port for a service other than the one intended.
- You can use intrusion detection system tools in a simulated network traffic or live network traffic test to aid you in determining if your packet filtering rules are protecting your systems and networks from known attacks. You will need to run these tools for some period of time and review the results on a regular basis.

You may want to defer this level of testing to normal operations once you have fully deployed the new firewall system.

- Examine all of the network traffic that is logged and verify that the logging options associated with each packet filtering rule are operating as intended.
- Examine all of the network traffic that is logged and verify that the alert options associated with each logging option are sending alerts to the designated destination (such as the firewall administrator) using the specified mechanism (such as paging or email).

You cannot do a final test of your routing configuration prior to connecting the firewall system to your operational external interfaces (refer to “9. Install the firewall system.” and “10. Phase the firewall system into operation.”). As a result, you should run live packets through your internal networks using the new firewall system to the greatest extent possible prior to connecting to the outside world. To mitigate the risk of unexpected problems in this final test phase, you should initiate the operational connections for a small subset of hosts (such as those used by your system and firewall administrators) prior to connecting large numbers of user workstation or server hosts.

➤ *Select and test features related to log files.*

When log files are full, you need to select how the firewall system will respond. The possible options may be one or more of the following

- shut down all external interfaces connected to the firewall
- continue to operate, overwriting the oldest entries
- continue to operate without logging

The first option is the preferred one but may not be available with all firewall products. Simulate a firewall log full condition and ensure that the firewall system behaves as expected based on the option you selected.

Select and exercise the appropriate settings for the archival of log files. The settings may include

- log file destination (e.g., a local file on the firewall host or a central log file on a remote host)
- number of days before archiving a specified log file
- number of days before purging an archived version of a specified log file

➤ *Test the firewall system*

For each relevant failure mode described in the test plan (see the first step), execute the test scenario causing or simulating the failure and exercise the mitigation strategy to see that it has the desired effect.

➤ *Scan for vulnerabilities*

Use vulnerability detection tools to scan your firewall system to determine the presence of known vulnerabilities. If patches exist for vulnerabilities that a tool detects, install these on your firewall system and re-execute the tool. This ensures that the vulnerability has been eliminated.

➤ *Design initial regression testing suite*

Select a subset of test cases to be used for regression testing purposes during normal operations. These should include cases that verify that all incoming and outgoing packets are being routed, filtered, and logged as expected as well as service-specific cases that verify that packets requesting specific services (WWW, email, FTP, etc.) are being routed, filtered, and logged as expected.

Once the new firewall system is part of normal operations, you can use selected regression test cases to verify that a change does not affect operational capabilities that worked successfully prior to the change.

➤ *Prepare system for production use*

Create and record cryptographic checksums or other integrity-checking baseline information of your firewall system once you have completed testing. Refer to *Preparing to Detect Signs of Intrusion* [Kochmar 98], specifically the practice "Generate information required to verify the integrity of your systems and data."

Make a backup of your operational configuration once you have completed testing. Refer to *Securing Desktop Workstations* [Simmel 99], specifically the practice "Configure computers for file backups."

➤ *Prepare to perform ongoing monitoring*

Given the complex nature of networks, their traffic, and firewall systems, ongoing monitoring is the only way to ensure that you have specified the correct security policy and that the policy is being implemented properly.

Ensure that you have the necessary policies, procedures, tools, and staff resources in place to monitor your networks and systems including your firewall system.

Policy considerations

Your organization's networked systems security policy should require that

- testing the firewall system is performed in an environment isolated from your operational networks
- the firewall system be retested after every configuration change and periodically using the regression test suite
- the regression test suite be kept up to date to exercise the current firewall system configuration
- the inventory of all applications software, operating systems, supporting tools, and hardware be kept up to date
- monitoring of all network and systems, including your firewall system, be performed on a regular basis

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p060.html>

Figure 8-1: Test Environment




Figure 8-2: Single layer firewall architecture




Figure 8-3: Multiple layer firewall architecture





Figure 8-4: Production environment



9

Install the firewall system.

A tested firewall system needs to be installed in the production environment given that there are generally differences between the test and production environments.

There are two cases to consider:

- The firewall system is being installed to provide new connectivity between the networks that it was designed to interconnect.
- The firewall system is being installed as a replacement for an existing system that already provides connectivity between these networks.

The approaches for these two cases differ and are described below.

Why this is important

Planning the installation and executing a defined plan are necessary to minimize interruptions in existing services.

How to do it

➤ *Install new connectivity*

Do not allow unfiltered traffic through the firewall system during installation. Turn off the firewall system, reinstall it, bring it back up carefully, and perform sample tests to prove it is forwarding and filtering as expected.

Make sure that the networks are prepared to be interconnected. Take into account, for example, IP addressing, routing, and DNS. Refer to *Firewalls and Internet Security* [Cheswick 94], *Building Internet Firewalls* [Chapman 95], and *Firewalls Complete* [Goncalves 98] for additional information.

Consider making any new services available incrementally (a few at a time). The easiest way to accomplish this is to insert a “deny all” filter into each rule set immediately after the services you wish to make available. To make more services available, move the “deny all” rule further down in the rule sets until it gets to the bottom and then remove it (it should be redundant at that point).

➤ *Install replacement connectivity*

Install the new firewall system in parallel with the existing system. Inserting the new firewall system into your production environment should not cause any changes to the environment, if possible.

Do not allow unfiltered traffic through the new firewall system during installation. Turn off the firewall system, reinstall it, bring it back up carefully, and perform sample tests to prove it is forwarding and filtering as expected.

During initial installation, maintain the existing system once it is disconnected. You can then switch back to it if the new system does not operate properly.

Policy considerations

Your organization's networked systems security policy should require that the firewall system installation/deployment plan and schedule are consistent with your site infrastructure business plan and schedule of infrastructure upgrades.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p061.html>

10

Phase the firewall system into operation.

Once you have physically installed your tested firewall system into your production environment, you can then integrate it into your operational networks. This practice addresses the case where your firewall system is replacing an existing firewall or router. In this configuration, no hosts are aware of the newly installed firewall. Each host that is intended to send traffic through the firewall must be made aware of the new firewall's existence. You then make sure that the packet filters perform as expected in the production environment.

There are two subcases to consider

- the replacement firewall system is installed using the same IP addresses as the original system
- the replacement firewall system is installed using different IP addresses than the original system

This practice includes steps for which the new, replacement firewall system is being installed using different IP addresses as this is the more complex case.

Why this is important

The hosts on all of the networks that your firewall system controls must be configured to send traffic to and receive traffic from the new firewall system or the firewall will not serve its purpose.

How to do it

- *Prepare for transition to the replacement firewall system*

Plan how to update all hosts connected to the new, replacement firewall with the routing information they each need to route traffic through this firewall. There are generally three ways to accomplish this:

- Manually configure the routing information on each host.
- Use one or more routing management protocols (e.g., RIP¹, OSPF²) to transmit routing information from the firewall to hosts.
- Use a client management protocol (e.g., DHCP³) to reconfigure hosts from a central server.

-
1. routing information protocol
 2. open shortest path first
 3. dynamic host configuration protocol

DHCP is supported by most TCP/IP stacks provided by operating systems along with a DHCP server that responds to requests for IP addresses. You can configure your DHCP server to provide default route addresses to clients. Change the default route for specified hosts that are to send their traffic to and receive their traffic from the new firewall.

Alternately, you may choose to distribute routing information to your client hosts via a routing protocol. The two most common protocols for this purpose are RIP and OSPF. The choice between these two protocols is beyond the scope of this document, but most TCP/IP routing texts cover both in depth⁴. For our purposes here, they are equivalent: both can distribute a default route to hosts that are listening for it.

In the absence of these capabilities (DHCP, RIP or OSPF), you will need to reconfigure each host to manually change the host's default route address.

You should maintain a fallback configuration to continue operations if the firewall system does not work as intended. Otherwise, you run the risk of incurring network outages that may affect the ability of your organization to conduct business that relies on internal and external communication via networks and through the firewall system. You should plan to perform the transition during non-business hours such as over a weekend.

➤ *Notify users*

Alert your users that the firewall system is being brought into your operational environment.

Inform them that the default gateways on their hosts will be changed to route network traffic through the firewall system and that this should be invisible to them. Indicate that they should inform their system or firewall administrator if they encounter any problems.

➤ *Enable private traffic through the new firewall system*

This step assumes that you are migrating from a single layer firewall architecture (see figure 10-1 at the end of this section) to a multiple layer architecture (see figure 10-2 at the end of this section). The old firewall system becomes the interface to the external world and the new firewall system serves as the interface to your internal networks. This step also assumes that you have a network topology of one or more private networks and one or more public networks. The public networks typically connect hosts that respond to internal and external requests for service (such as WWW [HTTP], FTP, email [SMTP], and DNS). These hosts may also respond to internal requests for services such as SNMP and logging. The public network as described here can serve as your DMZ. The private network typically connects hosts that service your internal users including individual user workstations.

Configure and enable packets generated by the hosts on your private network to pass through the new firewall system:

1. Connect your public and private networks to the new firewall system. See figure 10-2 at the end of this section.
2. Change the default gateway of the hosts on the private network from the old firewall system to the new firewall system. See figure 10-3 at the end of this section.

4. Refer to the list of Protocol references in the Executive Summary of this module.

3. Update the routing table on all public network hosts to route private network traffic through the new firewall system (vs. the old firewall system). See figure 10-4 at the end of this section.
4. Disable the interface to the private network on the old firewall. Add a route to the old firewall system's routing table to route private network traffic through the new firewall system (see figure 10-5 at the end of this section).
5. Ensure that traffic is being routed and filtered as expected. Refer to "8. Test your firewall system."
6. Unplug the private network interface to the old firewall system.

If you are replacing an existing firewall system, maintain the physical connectivity through your existing system when you bring the new one online. This will allow you to determine if there are any hosts on your private network that have not had their routing configuration updated to interface with the new firewall system (as their traffic will continue to flow through the old firewall system). It will also allow you to make sure that everything is working as expected (i.e., all private network traffic is now being routed through the new firewall system).

If this does not work as intended:

1. Unplug the new firewall system hardware.
2. Plug the old firewall system hardware back into the private network.
3. Change the old firewall system's private network interface address to the one that was used by the new firewall system, given all of your private network hosts are now using it.
4. Add the route for the private network in the old firewall system's routing table.

Policy considerations

Your organization's networked systems security policy should

- require that your users are notified of firewall system service outages in advance
- specify who is to be notified in the event of firewall system anomalies once deployment is complete

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p063.html>

Figure 10-1: Single layer firewall architecture




Figure 10-2: Multiple layer firewall architecture




Figure 10-3: Change default gateway




Figure 10- 4: Update routing

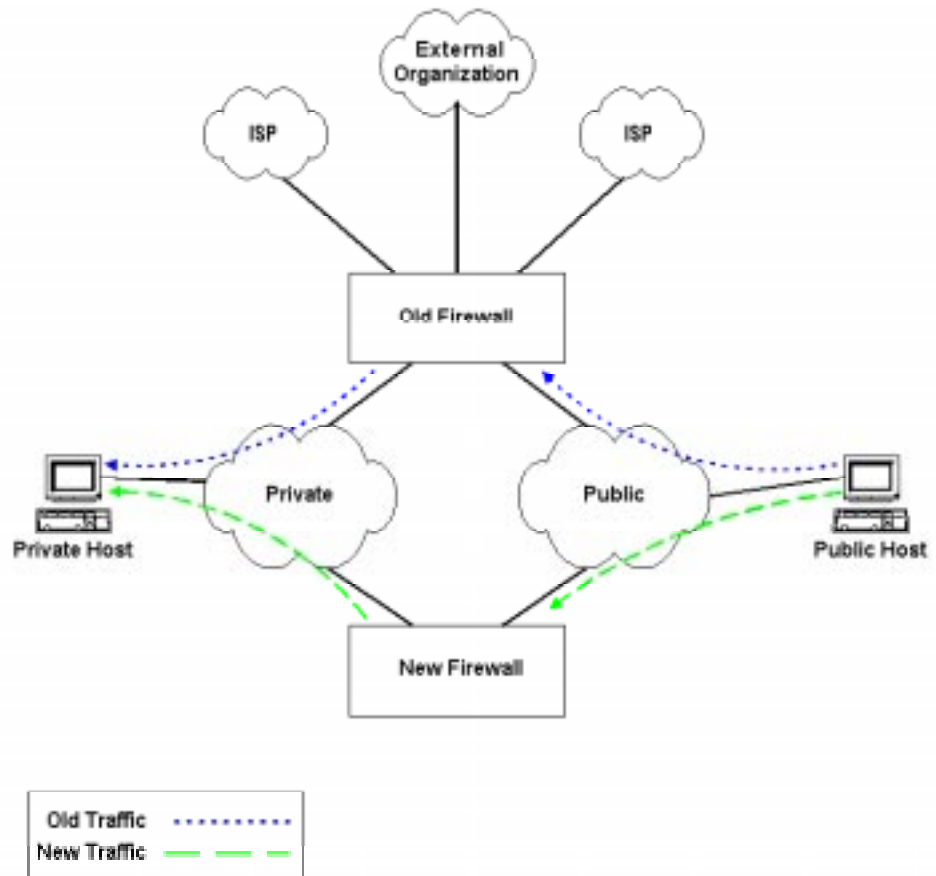


Figure 10-5: Disable old interface; add new route

